On the Applicability of “pgmcc” to UMTS Multicast

Antonios Alexiou[1, 2], Christos Bouras[1, 2], Andreas Papazois[2]
[1]*Research Academic Computer Technology Institute, Greece*
[2]*Computer Engineering and Informatics Dept., Univ. of Patras, Greece*
alexiua@cti.gr, bouras@cti.gr, papazois@ceid.upatras.gr

Abstract

In this paper, we investigate the applicability of the Pragmatic General Multicast Congestion Control (pgmcc) scheme to the multicast data delivery over Universal Mobile Telecommunications System (UMTS) networks. We show that the pgmcc scheme cannot support the degradation of the radio channels in the UMTS Terrestrial Radio-Access Network (UTRAN). Our major contribution is that the legacy scheme is modified in order to cope with the packet losses caused by either the temporary or the permanent degradation of the radio channel. The proposed scheme introduces minor modifications in the UMTS nodes and respects the limited computing power of the mobile equipment. Finally, we simulate our approach and we evaluate it under various conditions. The simulation results are presented along with their analysis.

1. Introduction

Third Generation (3G) mobile cellular networks promise the provision of real time multimedia services which address to user groups. UMTS constitutes the most prevalent standard of the 3G networks. Despite the high capacity that UMTS networks provide, the expected demand will certainly overcome the available resources. This is the reason why multicast transmission is one of the major goals for UMTS networks [4]. Multicast is an efficient method for transmission to multiple destinations. 3G Partnership Project (3GPP) recognized the need for the support of multicast in UMTS networks. As a result, the standardization of Multimedia Broadcast/Multicast Service (MBMS) framework started [1].

Congestion control is a policy that regulates the source transmission rate according to the network congestion. In IP multicast no congestion control is implemented. Instead, the Transmission Control Protocol (TCP) regulates its transmission rate according to network congestion. This means that the coexistence of multicast traffic and TCP traffic may lead to unfair use of network resources. In order to prevent this situation, the deployment of multicast congestion control is indispensable. This kind of congestion control is well-known as TCP-friendliness.

In this paper, we investigate the applicability of the well known pgmcc scheme [5] for the multicast congestion control over UMTS cellular networks. The adoption of a multicast congestion control in cellular networks poses an additional set of challenges which are related to the existence of wireless links. We show that the legacy pgmcc scheme cannot support the degradation of the radio channels in UTRAN. The innovation of our work stems from the fact that pgmcc is partly modified and extended in order to support the particularities of UTRAN. The proposed scheme introduces minor modifications in the UMTS nodes. Last but not least, our scheme respects the limited computing power of the mobile equipment.

Multicast congestion control problem in fixed networks is still a domain of active research and a lot of solutions have been proposed until now. In contrast to the multicast congestion control problem in fixed networks, very few solutions and algorithms have been proposed for the variation of this problem in cellular networks. The most strongly related publication is our work presented in [3]. In that paper we presented a novel mechanism based on another well known congestion control algorithm, the TCP-Friendly Multicast Congestion Control (TFMCC). TFMCC is an equation-based congestion control algorithm which differs considerably in the smoothness and the predictability of its transmission [7].

Another related work is presented in [6]. In this publication, the wireless-caused representative selection fluctuation problem is presented. This situation causes frequent change of the representative in wireless multicast congestion control. The sender adjusts its transmission rate to the tentative worst receiver, which brings severe performance degradation to wireless multicast.

The paper is structured as follows: Section 2 provides an overview of the UMTS architecture. In Section 3 we briefly present the pgmcc. The problem of the applicability of pgmcc over UMTS networks is described. Section 4 is dedicated to the proposed scheme and Section 5 describes the simulation experiments. Finally, some concluding remarks and planned next steps are stated in Section 6.

2. Overview of UMTS Architecture

From the physical point of view, the UMTS architecture is organized in two domains: the User Equipment (UE) and Public Land Mobile Network (PLMN). The UE is used by the subscribers to access the UMTS services while the PLMN is a network established by an operator to provide mobile services to the public. The PLMN is further divided into two land-based infrastructures: the UTRAN and the Core Network (CN) (Figure 1). The UTRAN handles all radio-related functionalities. The CN is responsible for maintaining subscriber data and for switching voice and data connections.

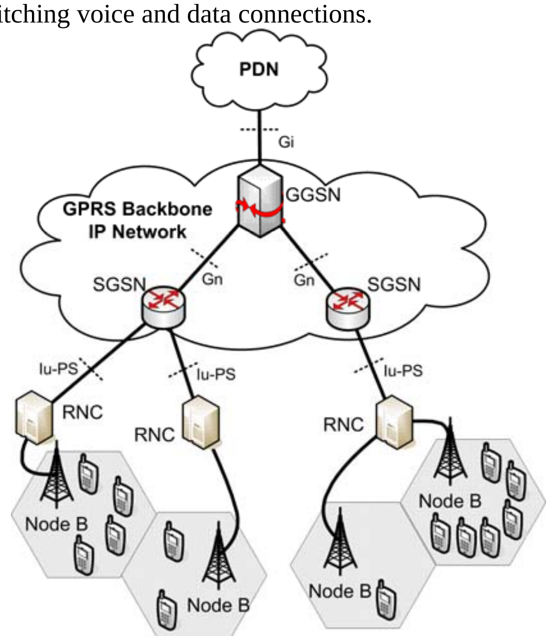

Figure 1. UMTS architecture.

The UTRAN consists of two kinds of nodes: the Radio Network Controller (RNC) and the Node B. The Node B constitutes the base station and provides radio coverage to one or more cells. The Packet Switched (PS)-domain of the CN consists of two kinds of General Packet Radio Service (GPRS) Support Nodes (GSNs), namely Gateway GSN (GGSN) and Serving GSN (SGSN). The SGSN provides routing functionality and manages a group of RNCs. GGSN provides the interconnection between the UMTS network and the external data networks [4].

3. "pgmcc" for Congestion Control

3.1. The"pgmcc" Scheme

In general, the receiver reports are a fundamental component of pgmcc. They are sent back to the sender as NAKor ACK fields and based on them the sender regulates the transmission rate. The NAKs consist of the following fields:

- the identity of the receiver (`recv_id`),
- the loss rate measured locally (`recv_loss`),
- the highest data packet sequence number received (`recv_lastseq`).

Additionally, the receiver with the worst throughput is elected as acker being in charge of sending positive ACKs to the sender. The identity of the acker is carried as a field in each data packet. The ACKs contain the same loss report as NAKs and two additional fields:

- the sequence number of the data packet which triggered the ACK (`ack_seq`),
- a 32-bit bitmap indicating the receive status of the 32 last packets (`bitmask`).

The loss rate is estimated in each receiver. The estimated loss rate (`recv_loss`) is sent back to the multicast sender with NAKs and ACKs. In order to measure its loss rate, each receiver interprets the packet arrival rate as a discrete signal (1 for lost packets, 0 otherwise). The signal is then passed through a discrete-time linear filter, whose response and computational costs are chosen accordingly. The sender regulates the transmission rate based on the worst throughput received.

The main goals of pgmcc are TCP-friendliness and scalability. The latter is achieved by decentralizing functionality as much as possible. The adoption of NAK suppression and the election of a single node to be in charge of sending ACKs are solutions which only require a constant amount of processing on nodes, independently of the group size.

Another advantage is fast response which is achieved by introducing positive ACKs from the acker instead of NAKs proposed by other schemes. The adoption of positive ACKs permits a more timely distribution of information.

3.2. "pgmcc" for UMTS

The applicability of pgmcc for the multicast congestion control over UMTS, faces a major problem. In cellular networks, there are some cases where the packet loss may not mean network congestion. Actually, the quality of a radio channel may be degraded and the bit error rate of the wireless link may become very high but, normally, after that period the wireless link is expected to recover.

In case a packet loss occurs, the legacy pgmcc translates the packet loss as buffer overflow in the network nodes, i.e. as network congestion. Consequently, the action taken in order to resolve this situation is the reduction of the sender's transmission rate. Nevertheless, if the packet loss is caused by radio channel degradation, the reduction of the transmission rate will not affect the packet loss. This is due to the fact that the packet loss does not depend on the arrival rate of the packets but on the radio channel degradation. Finally, the packet loss will be resolved after the end of the fading period, without a transmission rate regulation.

Obviously, the radio channel degradation may affect the performance of the pgmcc scheme. If we suppose that a UE suffers from fading, then the increment of packet loss may cause the election of this UE as the acker. The next step is that the transmission rate of the multicast server will be reduced according to the packet loss of the examined UE. After the recovering from bad wireless quality phase, the target throughput of the acker will be improved. Soon, another UE suffering from channel degradation will be selected as acker and will regulate the transmission rate. Eventually, the radio channel degradation will cause a significant and steady degradation of the performance of the pgmcc mechanism and of the multicast service. As a consequence, the network reaction to the packet losses due to radio channel degradation should not be a drastic reduction of the sender's transmission rate. During this analysis we shall refer to this problem as the acker election problem.

4. The Modified Scheme

4.1. Modification of Packet Formats

In order to cope with the acker election problem, the original pgmcc scheme was modified. Figure 2 illustrates the packet formats used by the modified pgmcc scheme. The dark grey field is new, whereas the light grey fields are modified to support the new scheme.

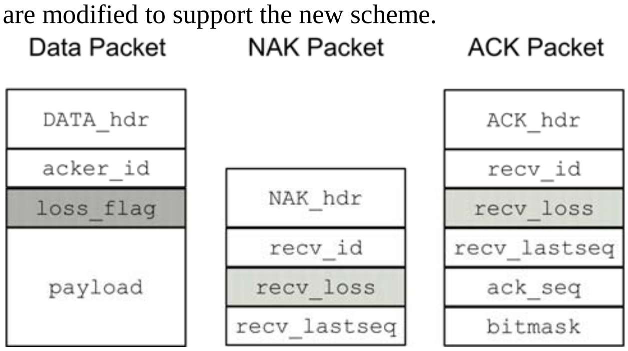

Figure 2. Packet formats for modified scheme.

A new packet field has been added in the data packet of the proposed scheme. This new field is described by the identifier `loss_flag` and its usage will be described in the following subsection.

On the contrary, the same formats are preserved in the modified scheme for the NAK and the ACK packets. Nevertheless, the way the `recv_loss` field is calculated has changed.

The way that receiver loss rate field is calculated depends on the type of the radio channel degradation. The radio channel degradation may be temporary or permanent. Although the effects of temporary radio channel degradation should be filtered out by the scheme, the permanent radio channel degradation impacts should not be skipped. In the following subsections we explain these situations and we describe thoroughly the way that the receiver loss rate field is calculated under these circumstances.

4.2. Temporary Radio Channel Degradation

By the term temporary radio channel degradation we mean any situation which causes increment of the packet losses over the radio channel for a short time period. The main goal of the modified scheme is to filter out the packet losses caused by this situation. The reason is that this kind of packet loss does not mean network congestion and, consequently, no reduction of the sender's transmission rate is needed.

In the proposed mechanism, the nodes located at the border between wireless and wired network (i.e. the Node Bs) have an additional responsibility. This responsibility is to provide the receivers (i.e. the UEs) with information about their measured packet loss. This information is piggybacked in the data packets of the multicast session. It is written as new field of the data packets and is then read by the UEs. This is a totally new functionality introduced to combat the acker election problem in UMTS networks.

This additional functionality of the Node Bs, permits each UE to identify the reason of a packet loss. In general, the following cases are distinguished:

- When both the Node B and the UE encounter a packet loss, this packet loss is considered to be caused due to network congestion. Consequently, this kind of packet loss is taken into consideration during the acker election.
- On the other hand, when the two values differ, the UE can conclude that the reason for the difference is losses at the wireless link caused by radio channel degradation. In this case, the packet loss is not accounted during the acker election.

The packet loss flag field introduced in the data packet of the modified scheme (`loss_flag` in Figure 2) is set by the Node B and may be an 1-bit information. This means that 0 indicates a correctly received packet in Node B, whereas 1 indicates a previously lost packet. Each UE reads this information when receives a data packet. The `loss_flag` sequence from the data packets is recorded in the UE and is used for the calculation of the loss rate via the discrete-time filter.

The intention of the modified scheme is to exclude the packet losses caused by radio channel degradation from the input signal of the discrete-time filter. Obviously, the desired sequence coincides with the loss flag sequence indicated by the Node B. The reason is that if a packet

loss occurs in the Node B then the same packet will definitely not reach the connected UEs as well.
The packet loss flag sequence is used as input for the discrete-time filter of the legacy pgmcc scheme. Figure 3 visualizes the basic principle of the operation.

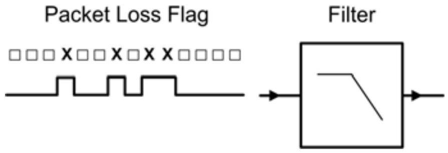

Figure 3. The loss flag indication is used as filter input.

The output of the filter is the receiver's loss rate estimation which is sent with the feedback packets (ACKs or NAKs) to the multicast sender. Based on this feedback, the multicast sender makes the acker election.

4.3. Permanent Radio Channel Degradation

Another aspect of the proposed mechanism is the case that, under certain conditions, a permanent degradation of the radio channel affects a specific UE. When a permanent degradation occurs on the wireless link, the buffer of the Node B will overflow and some packets will be rejected. Normally, this UE should be an acker-candidate.
In the proposed mechanism, during permanent channel degradation, the Node B counts the rejected packets as general packet losses which happened due to network congestion. The packet loss flag for these packets is set by the Node B and, consequently, these packets are taken into account by the UE during the calculation of the loss rate. Figure 4 illustrates this operation of the Node B. The packet loss flag is set if either the packet is lost or it is rejected from the buffer of the Node B due to permanent radio channel degradation.

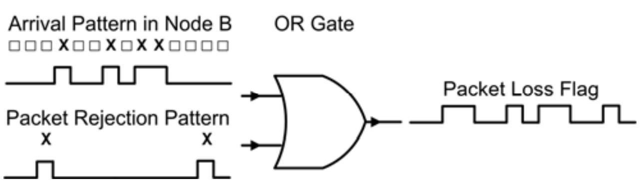

Figure 4. Determination packet loss flag in Node B.

As we mentioned above, it is assured that the permanent degradation is not hidden from the UE. In fact, the UE is informed of the packet losses caused by buffer overflow. This functionality makes our proposed mechanism suitable not only with the temporary, but also with the permanent degradation of the radio channel.

5. Experimental Results

The above described mechanism was evaluated towards two directions. The first was to examine that the proposed mechanism preserves the benefits of pgmcc as they are presented in [5]. The other was to evaluate the behavior of the modified scheme against the acker election problem when temporary or permanent radio channel degradation occurs.
For the verification of the proposed scheme the ns-2 network simulator was used along with the multicast packet forwarding mechanism described in [2]. Additionally, the implementation of the legacy pgmcc, provided by the authors of [5], was enhanced in order to support the functionality of the Node B and the UE as described in Section 4.

5.1. Intra-Protocol Fairness

The first aspect that we examined was the intra-protocol fairness of the proposed scheme. In other words, we considered the fairness of the proposed scheme towards other competing pgmcc flows when they share common wired or wireless links.

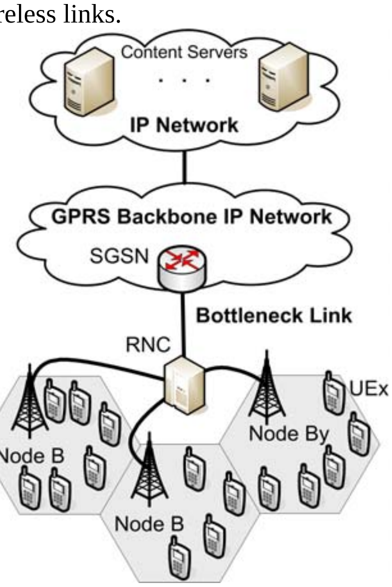

Figure 5. Single-bottleneck topology.

In the first place, we simulated the congested UMTS network depicted in Figure 5. It is a single-bottleneck topology with a bottleneck applied over a link which connects an SGSN with an RNC node (Iu interface). We monitored the throughput over a wireless link connected the UEx with Node By. We supposed that UEx belongs to two multicast groups and receives two instances of pgmcc traffic from two different external nodes.
Figure 6 illustrates the throughput of each competing pgmcc flow.
As it was expected, the available bandwidth is fairly shared between the flows. Figure 6 confirms that the

average throughput of pgmcc1 flow closely matches the average pgmcc2 throughput.

Similar results were obtained for many other congestion scenarios. For example, if we suppose that no congestion exists or that congestion exists over a Gn interface (connects GGSN with SGSN nodes). The intra-protocol fairness of the proposed scheme was therefore confirmed under all the congestion scenarios.

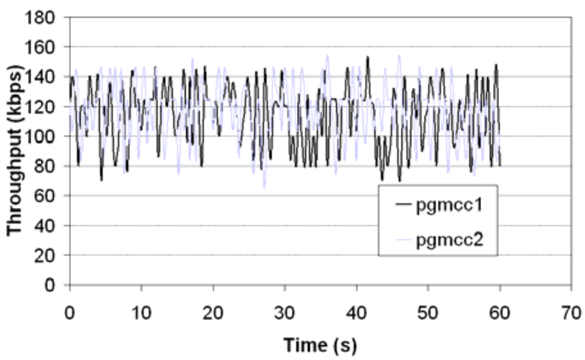

Figure 6. Fairness between two pgmcc instances.

5.2. Inter-Protocol Fairness

The inter-protocol fairness of the proposed scheme is the fairness towards the competing TCP flows when they share wired or wireless links. Obviously, this concept is identical to the concept of TCP-friendliness. In order to examine the inter-protocol fairness of the proposed scheme, the same single-bottleneck UMTS topology was used. We supposed that UEx belongs to a multicast group and receives pgmcc traffic. At the same time, this UE receives TCP traffic from an external node.

We monitored the throughput of a pgmcc flow against some sample TCP flows. The outcome was that the average throughput of pgmcc closely matches the average TCP throughput. Moreover, pgmcc achieves a smoother rate than the TCP.

In all examined scenarios, the available throughput of the bottleneck link is evenly shared among the competing pgmcc and TCP flows. The inter-protocol fairness of the proposed scheme was therefore confirmed.

5.3. Responsiveness to Changes

An important concern in the design of congestion control protocols is their responsiveness to changes in the network conditions. This behavior was investigated using the bottleneck topology of Figure 5. During the simulation we applied three different loss rates on the bottleneck link. The pgmcc flow was monitored along with two TCP flows sharing the bottleneck link. The results of the simulation for the three competing flows are presented in Figure 7. The proposed scheme throughput matches closely the TCP throughput at all three loss levels. Moreover, the adaptation of the sending rate is fast enough.

A similar simulation setting was used in order to investigate the responsiveness to changes in the RTT. The results are similar to those above. The above experiment confirms the excellent reactivity of the pgmcc to changes in congestion level of the UMTS network.

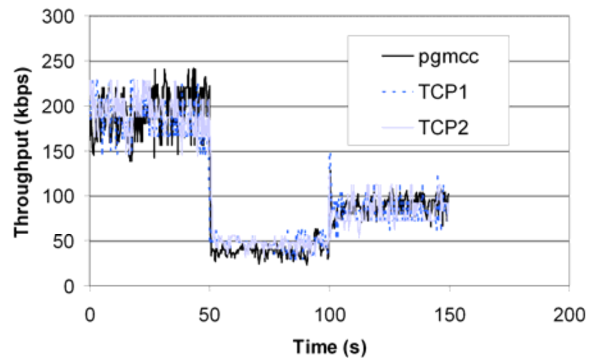

Figure 7. Responsiveness to changes in the loss rate.

5.4. Temporary Radio Channel Degradation

The next concern of our experiments was the evaluation of the proposed scheme when packet losses occur due to temporary radio channel degradation. We simulated a UMTS network and assumed degradation on the radio channels by applying an error rate over the packets transmitted. We examined the proposed scheme for different number of UEs belonging in the multicast group.

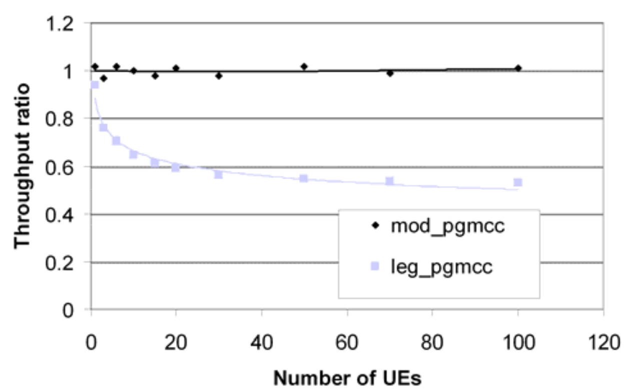

Figure 8. Throughput of proposed scheme vs. legacy pgmcc during temporary radio channel degradation.

In Figure 8, our proposed scheme is referred as "mod_pgmcc" (modified pgmcc). On the other hand, the legacy pgmcc algorithm is referred as "leg_pgmcc". The horizontal axis shows the number of the UEs belonging in the examined multicast group. Both mechanisms were examined for up to 100 UEs participating in the multicast

group. The vertical axis shows the average throughput which is normalized to the corresponding TCP one. The results when 5% packet loss is applied are presented.
The results depicted in Figure 8 confirm that the packet losses can be identified correctly at the UEs and be ignored at the calculation of the acceptable sending rate. This means that the acker election problem can be overcome and significant improvement is added on the pgmcc application over the UMTS.

5.5. Permanent Radio Channel Degradation

The last concern of our experiments was the evaluation of the proposed scheme when packet losses occur permanently due to radio channel degradation. We simulated the degradation by applying an error rate of 30% over the packets transmitted via the corrupted wireless link. After 50 seconds of simulation, we applied the error rate over the wireless channel connecting the UEx with Node By. We monitored the changes over the throughput of the corrupted wireless link for 100 seconds. The results of our experiment are presented in Figure 9.

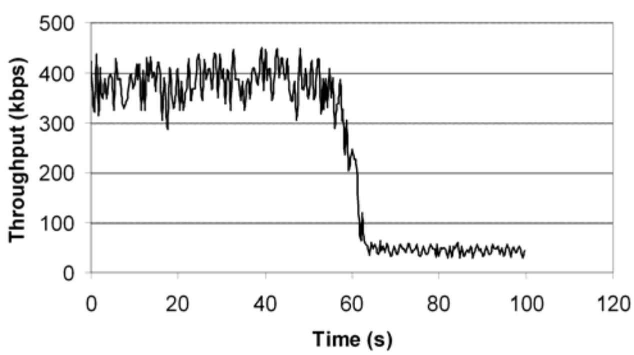

Figure 9. Throughput of proposed scheme during permanent radio channel degradation.

The simulation results prove that our scheme reacts to the permanent wireless-channel degradation. In the beginning of the simulation, no congestion exists and the throughput matches with the available bandwidth of the wireless link. When the 30% packet error rate is applied, the network does not immediately react to this degradation because it considers it as wireless-caused degradation.
During the period between 50 and 60 seconds, the buffer of the Node B overflowed and the network was able to distinguish the nature of the degradation. It took about 10 seconds to adapt to the new network conditions, but this time interval may differ according to the bit-rate of the transmission and the size of the buffer in Node B.
The results confirm that the packet losses in the Node B are considered as network congestion and are not ignored during the calculation of the acceptable sending rate in the UE. Eventually, this kind of packet losses causes reduction of the transmission rate.

6. Conclusions and Future Work

We have concluded pgmcc had to be modified before being applied to UMTS networks. The pgmcc scheme was modified so as not to translate the radio channel-caused packet losses as buffer overflow in the network.
We have evaluated the proposed scheme through simulation experiments. We concluded that it preserves the benefits of pgmcc algorithm i.e. intra and inter-protocol fairness and responsiveness to changes. In addition, the results confirm the excellent behavior of our proposed scheme when temporary or permanent radio channel degradation occurs. Minor modification in the UMTS architecture is needed in order to adopt the new scheme. Actually, the impacts concern only two nodes of the UMTS network; the Node Bs and the UEs.
The step that follows this work may be the formulation of a multicast group control mechanism dedicated for the UMTS networks. It may be specified in which cases the permanent radio channel degradation may cause a large reduction to the transmission rate and eventually multicast service deterioration.

7. References

[1] 3GPP, *Multimedia Broadcast/Multicast Service (MBMS); Architecture and functional description (Release 7)*, Technical Specification 23.246 v7.2.0, March 2007.
[2] A. Alexiou, D. Antonellis, C. Bouras, and A. Papazois, “An efficient multicast packet delivery scheme for UMTS”, *International Workshop on Modeling Analysis and Simulation of Wireless and Mobile Systems (MSWiM’06)*, ACM Press New York, NY, USA, 2006, pp. 147-150.
[3] A. Alexiou, C. Bouras, and A. Papazois, “Multicast in UMTS: Adopting TCP-Friendliness”, *Personal Wireless Communications (PWC 2007)*, Springer, 2007, pp. 433-444.
[4] H. Holma, and A. Toskala, *WCDMA for UMTS: Radio Access for Third Generation Mobile Communications (3rd Edition)*, Wiley, 2004.
[5] L. Rizzo, “pgmcc: a TCP-friendly single-rate multicast congestion control scheme”, *ACM SIGCOMM Computer Communication Review*, 30, 4, ACM Press New York, NY, USA, 2000, pp. 17-28.
[6] T. Saito, and M. Yamamoto, “Wireless-caused representative selection fluctuation problem in wireless multicast congestion control”, *IECE Transactions on Communications*, E88-B, 7, Oxford University Press, Oxford, 2005, pp. 2819-2825.
[7] J. Widmer, and M. Handley, “Extending equation-based congestion control to multicast Applications”, *ACM SIGCOMM Computer Communication Review*, 31, 4, ACM Press New York, NY, USA, 2001, pp. 275-285.